The 10-Bit Personality: A Binary Audience-Fingerprint Across Partisanship, Israel–Palestine, Secularism, Church Authority, Charismatic Practice, Conspiracy Frames, AI Futures, Academic Gatekeeping, and Nationalism

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

January 17, 2026

This paper uses “10-bit” in its strict information-theoretic sense: ten binary choices yield 2^10 = 1024 possible signatures. On the face of it, 1024 cannot uniquely identify a person in a world of billions. But identifiability is always relative to a reference set. Among public-facing writers and public figures, two conditions change the math. First, they are repeatedly forced into yes-or-no outputs by media formats, coalition loyalty tests, and the economics of attention. Second, most ordinary people do not hold crisp opinions across the same ten contested domains, or they can remain silent without penalty. Public figures often cannot. Here we define ten specific “bits” as the presence or absence of friction with ten audience blocs that routinely demand binary conformity: hard partisans, maximalists on both sides of Israel–Palestine, strong secularists, church institutional loyalists, charismatic communities keyed to tongues practice, conspiracy-curious and anti-conspiracy skeptics, AI accelerationists and AI doom or safety absolutists, academic gatekeepers and anti-academic populists, and nationalist audiences. The thesis is not that a 10-bit code captures a whole person, but that it can function as a practical stance-fingerprint in public life: compressive, fallible, yet surprisingly discriminative when most people are un-coded on most bits. The paper ends with coding rules, collision analysis, and ethical constraints for using stance fingerprints without dehumanizing the humans they compress.

Keywords: 10-bit personality, binary stance vector, audience fingerprint, public figures, polarization, partisanship, Israel Palestine, secularism, church authority, charismatic tongues, conspiracy skepticism, AI accelerationism, AI safety, academic gatekeeping, anti-academic populism, nationalism, information theory, entropy, coalition signaling, narrative warfare, disinformation ecology.

A collaboration with GPT-5.2-Thinking. 51 pages. CC4.0.

Outline

1. Abstract
A compact summary of the 10-bit audience-fingerprint model, why it can be near-unique among public figures, and the main cautions.

2. Introduction
2.1 What "10-bit" means here
Ten binary choices, 1024 signatures, uniqueness relative to a population.
2.2 Why public figures are a special population
Opinion pressure, reduced tolerance for silence, repeated demand for binary outputs.

3. The Ten Bits as Audience-Pressure Axes
This section defines each bit as a coding question: does the figure reliably produce language that this audience bloc experiences as aligned and affirming, or as irritating and disloyal. The bit is about friction, not about moral rightness.

3.1 Hard partisans in US politics
Binary demand: team confirmation versus refusal of team logic.
3.2 Pro-Israel maximalists and Zionist advocacy circles
Binary demand: unconditional legitimating frame versus sustained critique of policy and power.
3.3 Pro-Palestinian maximalists and anti-Israel absolutists
Binary demand: totalizing condemnation frame versus refusal of zero-sum moral narratives.
3.4 Strong secularists and anti-religious public-frame critics
Binary demand: strictly material public language versus theological or Logos-framed analysis.
3.5 Church traditionalists and institutional loyalists
Binary demand: deference to institutional authority versus prophetic critique of leadership and decline.
3.6 Charismatic and Pentecostal communities focused on tongues practice
Binary demand: affirmation of tongues as normative sign versus critique of uninterpreted practice and excess.
3.7 Conspiracy-curious communities and anti-conspiracy skeptics
Binary demand: everything-is-a-psyop certainty versus disciplined uncertainty and evidentiary restraint.
3.8 AI accelerationists and AI doom or safety absolutists
Binary demand: rapid deployment is salvation versus strong brakes as moral necessity.
3.9 Academic gatekeepers and anti-academic populists
Binary demand: strict disciplinary policing versus public-intellectual insurgency and anti-credential rhetoric.

3.10 Nationalist audiences across multiple countries

Binary demand: my-nation moral exceptionalism versus universal moral language that critiques all sides.

4. Why Ten Bits Can Feel Identifying

4.1 The missingness effect

Most people are un-coded on many bits because they do not speak on them. Public figures are more fully coded.

4.2 The repeat-exposure effect

Frequent statements reduce ambiguity and force a stable signature.

4.3 The salience effect

High-conflict topics generate strong audience reactions that sharpen the coding.

5. Why Ten Bits Often Are Not Identifying

5.1 Correlated bits and ideological bundling

If several bits move together, the effective information is less than ten.

5.2 Coalition templates and mass-produced personas

Many public figures are manufactured as predictable stance bundles.

5.3 Collisions and look-alikes

Different biographies and motives can map to the same friction pattern.

6. Coding Method: Turning Public Output into a Bitstring

6.1 Data selection rules

What counts as evidence: speeches, essays, interviews, long-form statements.

6.2 Coding rules per bit

Thresholds for setting a bit, handling ambiguity, and avoiding single-quote overfitting.

6.3 Confidence scoring

A companion score per bit that tracks certainty rather than forcing false precision.

6.4 Time and drift

How signatures change across years and why recency windows matter.

7. Worked Examples Without Name-Claims

This section uses archetypes rather than real-person identification, to show how the method discriminates while avoiding doxxing-by-inference.

7.1 The partisan amplifier

High alignment with partisan audiences, low friction with coalition enforcers.

7.2 The prophetic institutional critic

High friction with institutional loyalists, mixed friction elsewhere.

7.3 The technologist with AI maximal commitments

Low friction with accelerationists or safety absolutists, predictable collision patterns.

7.4 The cross-pressure peacemaker
High friction across multiple maximalist blocs, high distinctiveness within public-figure sets.

8. Interpretation Limits
8.1 Positions versus reasons
Same bits can arise from incompatible ethical and theological frameworks.
8.2 Audience reaction versus author intent
Friction is relational; it can be generated by misunderstanding as well as substance.
8.3 The person exceeds the code
What the bitstring compresses away: virtues, temperament, private charity, and inner life.

9. Ethical Constraints and Misuse Risks
9.1 Compression and dehumanization
Why stance fingerprints tempt reductionism and caricature.
9.2 Reputation harm and automated stereotyping
Risks in using bitstrings as labels, targeting tools, or exclusion mechanisms.
9.3 A humane protocol
Rules of use: provisional coding, explicit uncertainty, refusal to equate a code with a soul.

10. Applications
10.1 Journalism and debate moderation
Clarifying coalitional pressures without rewarding performative binary outputs.
10.2 Computational social science and discourse mapping
How to measure effective bits, correlation, and drift.
10.3 Personal authorship and self-consistency audits
Using the ten bits as a mirror for coherence, not as a branding weapon.
10.4 AI-assisted sensemaking
How AI can help map signatures while keeping human judgment accountable.

11. Conclusion
Why 2^10 can be practically discriminative in the public-figure regime, why it still collides, and why the ethical burden is to keep the human remainder visible.

12. References
All citations and sources collected here only, with no in-line references in the section expansions that will follow.

Art

1. Abstract

This paper defines a “10-bit personality” as an audience-fingerprint: a binary vector indicating whether a public figure’s public output tends to generate sustained friction with, or affirmation from, ten high-pressure audience blocs. The ten blocs are: hard partisans in U.S. politics; pro-Israel maximalists; pro-Palestinian maximalists; strong secularist critics of religious framing; church traditionalists and institutional loyalists; charismatic and Pentecostal communities where tongues practice is a boundary marker; conspiracy-curious audiences and anti-conspiracy skeptics; AI accelerationists and AI doom or safety absolutists; academic gatekeepers and anti-academic populists; and nationalist audiences across multiple countries. Each “bit” is operationalized as a repeatable coding rule based on observable patterns of praise, backlash, and coalition enforcement across time, rather than on the author’s private intent.

Although 2^10 yields only 1024 signatures, the paper argues that near-uniqueness can emerge within the subpopulation of public figures because public life penalizes silence and ambiguity. Media incentives, loyalty tests, and platform dynamics force many public figures to speak across multiple contested domains, reducing missingness and stabilizing a distinctive stance profile. However, the paper also shows why uniqueness is often illusory: many bits are correlated by ideological bundling, causing the effective information content to be far less than ten and producing frequent “collisions” in which different people share the same fingerprint.

The paper’s contribution is a disciplined framework for constructing, interpreting, and ethically constraining stance fingerprints. It outlines methods for confidence scoring, drift across time, and collision analysis, and it warns against reductionism, caricature, and reputational misuse. The central claim is modest but practical: a 10-bit audience-fingerprint can be a useful map of public cross-pressures, provided it is treated as provisional compression rather than as a total description of a person.

2. Introduction

The phrase “10-bit personality” is easily misheard as a celebration of nuance, emotional range, or high-resolution interiority. That meaning can be valid in other contexts, but it is not the meaning used here. This paper uses “10-bit” as a deliberately austere metaphor borrowed from information representation: a person’s public stance footprint can be approximated as a vector of binary choices across a set of high-pressure topics. The result is not a full portrait of a human being. It is a compressed signature of how a public communicator repeatedly collides with, or aligns with, predictable audience blocs that demand conformity.

The point of the model is not to trap a person in a label. The point is to describe a real feature of the present information ecology: public discourse increasingly behaves like a hashing function.

As soon as a writer becomes legible, coalition pressures begin to treat their words as a set of membership tests. The world does not read in paragraphs; it reads in bits. The discipline, then, is to study the bitmaking process rather than pretend it is not happening.

At the same time, the model must be handled with care. A stance fingerprint can clarify cross-pressures and ideological bundling, but it can also create caricature. It can be used to understand how a person is being interpreted by different crowds, or it can be used as a shortcut for condemnation. This paper argues that the model is valuable precisely because it is compressive and therefore dangerous: its usefulness depends on the explicit recognition that a human being always exceeds the code.

2.1 What "10-bit" means here

In a strict combinatorial sense, ten binary choices produce 2^10 distinct patterns, which equals 1024. If we imagine a set of ten yes-or-no questions and assign 1 for "yes" and 0 for "no," then each person's answers can be written as a 10-digit bitstring. This is the simplest possible stance representation: it records direction without recording intensity, reasons, or nuance.

It is immediately obvious that 1024 patterns cannot uniquely label every human being. The move that makes the model interesting is to treat identifiability as relative. A fingerprint is always defined with respect to a reference set. A handprint does not identify you among all mammals; it identifies you within a population where that kind of print is comparable and where the measurements are actually taken. In the same way, a 10-bit stance vector is not claimed to uniquely identify a person in general. It is claimed to be potentially discriminative within populations that are both (a) heavily observed and (b) repeatedly forced to produce comparable outputs.

Two further clarifications matter. First, in real life these "choices" are not cleanly observed as inner beliefs. They are observed as public patterns of speech and reception. That means the bitstring is best understood as an interactional signature: the stable shape of friction and affirmation between a speaker and an audience ecosystem. Second, 10 bits is not a magic number. The number of bits is less important than the independence and observability of the axes. Ten strongly correlated questions can collapse into two or three effective bits. Conversely, ten well-chosen, partially independent pressures can behave like a surprisingly discriminative code within a narrow but influential class of speakers.

So "10-bit" here means: a compact stance fingerprint composed of ten binary audience-pressure axes, yielding up to 1024 possible patterns, with uniqueness evaluated only relative to a defined population and a defined observation window.

2.2 Why public figures are a special population

The 10-bit model becomes plausible only when we notice that public figures are not sampled like ordinary individuals. Most people are not required to speak across ten contentious domains. Most people can remain silent, remain vague, or decline to answer without losing their livelihood, their platform, or their community standing. In everyday life, “no opinion” is an acceptable state. In public life, “no opinion” is often treated as evasive, cowardly, or complicit. This changes the data regime.

Public figures occupy a different environment in at least four ways.

First, they are subject to opinion pressure. Interviews, debates, podcasts, and social platforms are structured as extraction machines for positions. They compress complex moral and political landscapes into prompts that demand a side. The incentives reward speed, clarity, and confidence, even when the underlying reality does not warrant it. If a public figure declines to compress, that refusal becomes part of their persona and is itself interpreted as a stance.

Second, public figures face reduced tolerance for silence. Silence is not neutral; it is interpreted. On many topics, not speaking is treated as a hidden bit: either you are afraid of your own side, or you are secretly on the other side, or you are protecting the institution, or you are signaling elite complicity. The upshot is that missing data is actively punished, which means public figures are more fully coded across more axes than ordinary individuals.

Third, public figures experience repeated demand for binary outputs. Even when a public figure offers nuance, audiences often re-encode it into a bit. One crowd hears a single qualifying clause and calls them disloyal. Another crowd hears a single moral condemnation and calls them captured. Over time, the public record accumulates: clips, quotes, essays, statements, refusals, hedges, and walk-backs. This repetition stabilizes a signature. A single comment can be ambiguous; a thousand comments become a pattern.

Fourth, coalition enforcement turns many of the ten axes into loyalty tests. On the topics chosen in this paper, there exist audience blocs that do not merely disagree; they interpret disagreement as moral failure or identity betrayal. These blocs police boundaries, reward alignment, and punish deviation. That policing creates a feedback loop: as a public figure is rewarded or punished, they either conform, double down, or attempt to triangulate. Each response further clarifies their stance fingerprint.

These conditions explain why 2^10 can be “not enough” in general, but “often enough” in a public-figure regime. The model does not claim perfect identification. It claims that within a sufficiently observed, sufficiently pressured population, and with sufficiently salient axes, a 10-bit audience-fingerprint can become practically distinctive. But the same conditions that make it distinctive also make it ethically fraught, because the ease of coding tempts observers to

replace understanding with labeling. This paper proceeds with that tension in full view: the fingerprint can map cross-pressures, but it must never be mistaken for the person.

3. The Ten Bits as Audience-Pressure Axes

A stance fingerprint is only as good as its axes. In this paper the axes are not "issues" in the abstract, but audience-pressure blocs that behave like binary compressors. Each bloc tends to demand a recognizable rhetorical signal that confirms membership, loyalty, or shared moral framing. When a public figure speaks, these blocs do not merely interpret content; they interpret allegiance.

The purpose of defining the ten bits as audience-pressure axes is to shift the unit of analysis from private belief to public interaction. We are not coding what a person "really thinks." We are coding whether their public language reliably lands as aligned and affirming to a bloc, or as irritating, disloyal, and suspect. The code is about friction, not about moral rightness. A figure can be morally correct and still generate friction. A figure can be morally wrong and still generate affirmation. The bitstring is not a halo or a condemnation. It is a map of recurrent cross-pressures.

Each bit is therefore posed as a coding question that can be answered from observable public evidence: repeated statements, consistent rhetorical choices, the pattern of backlash and praise, the need for clarifications or walk-backs, and the social locations of the audiences reacting. The coding also depends on time windows. A person may flip a bit after a crisis, a conversion, a new coalition alignment, or a new institutional role. The goal is not to freeze a person, but to describe a period of recognizable public stance.

A final note before the ten bits are defined: none of these axes is perfectly clean. Each contains internal diversity and sub-factions. But the point of a bit model is not perfect sociological taxonomy. It is to capture the binary demands that repeatedly emerge at scale. A "bit" represents a recurring pressure to choose a side, not a claim that the side is coherent or righteous.

3.1 Hard partisans in U.S. politics

Hard partisans are defined here as audiences for whom political identity is fused to team membership. The binary demand is team confirmation versus refusal of team logic. The core test is not policy detail; it is whether the speaker performs loyalty.

The alignment signal is predictable: consistent "our side good, their side bad" framing, selective outrage, strategic omissions, and the use of in-group language that flatters the tribe's self-image. Hard partisans reward simplified causal stories that assign blame cleanly and treat

compromise as contamination. They also reward rhetorical aggression directed outward, because outward aggression functions as a loyalty display.

Friction appears when a figure refuses the team frame. That refusal can look like symmetrical critique, insistence on standards that apply to both sides, refusal to speak in enemy caricatures, or a willingness to acknowledge legitimate concerns across the aisle. Friction also appears when a figure demands evidence that undermines a favored narrative, or when they correct misinformation circulating inside their own coalition. In a hard partisan ecology, correcting your own side is treated as betrayal; correcting the other side is treated as leadership.

Coding rule in plain terms: if the figure's language repeatedly signals allegiance and avoids cross-tribal symmetry, the partisan bit tends toward alignment. If the figure repeatedly breaks the loyalty script and absorbs backlash from their own side for doing so, the bit tends toward friction.

3.2 Pro-Israel maximalists and Zionist advocacy circles

This bloc is defined by a binary demand: an unconditional legitimating frame versus sustained critique of policy and power. The key feature is not simple support for Israel's existence or concern for Jewish safety, which many people share. The maximalist version treats persistent public critique of Israeli state policy, military conduct, or nationalist ideology as illegitimate or as a disguised attack on Jewish identity itself.

The alignment signal is a framing that foregrounds Israel's security claims, emphasizes existential threat narratives, and tends to interpret opposition primarily as antisemitism or as irrational hatred. It often includes strong rhetorical defense of state actions as necessary, and a preference for language that minimizes or relativizes Palestinian suffering, especially when that suffering is placed in a moral ledger against Israeli conduct.

Friction appears when a figure repeatedly uses moral language that constrains state power, treats Palestinian civilian suffering as morally weighty and not merely unfortunate, or challenges settlement policy, collective punishment, dehumanizing rhetoric, or maximalist territorial claims. Friction can also appear when a figure refuses certain rhetorical taboos: saying "this policy is wrong" without retreating into purely procedural critique, or insisting on international-law style constraints rather than purely security-centered reasoning.

Coding rule in plain terms: if a figure's language is reliably experienced by this bloc as legitimating and protective of the state's moral frame, the bit tends toward alignment. If the figure's language repeatedly triggers accusations of betrayal, antisemitism-adjacent suspicion, or "giving aid and comfort" narratives, the bit tends toward friction.

3.3 Pro-Palestinian maximalists and anti-Israel absolutists

This bloc is defined by the opposite binary demand: totalizing condemnation versus refusal of zero-sum moral narratives. Again, the axis is not simple compassion for Palestinians or criticism of Israeli policy. The maximalist version treats any refusal of comprehensive delegitimation as complicity.

The alignment signal is a rhetorical posture that casts the conflict as a pure morality play with a single legitimate subject and a single illegitimate oppressor. It often rejects language that differentiates between a government, a military, a population, and a religious or ethnic identity. It may demand specific formulas, slogans, and boundary phrases as membership tests, and it can treat nuance as laundering.

Friction appears when a figure insists on moral constraints that apply across the board, rejects dehumanization of Israelis or Jews, distinguishes policies from people, or allows tragic complexity to remain visible. Friction also appears when a figure refuses the demand to treat every event as confirmatory of a single grand story, or when they name violence by non-state actors as morally real rather than as rhetorically excusable.

Coding rule in plain terms: if a figure's language repeatedly satisfies the demand for totalizing condemnation and rejects plural moral claims, the bit tends toward alignment. If the figure repeatedly gets attacked for "both-sidesing," "centering," "softening," or "normalizing," the bit tends toward friction.

3.4 Strong secularists and anti-religious public-frame critics

This bloc is defined by a binary demand: strictly material public language versus theological or Logos-framed analysis. The demand is not simply for good evidence or for separation of church and state. The strong version treats religious vocabulary in public reasoning as intrinsically irrational, manipulative, or disqualifying.

The alignment signal is a rhetorical style that keeps causality and meaning within a naturalistic frame: incentives, institutions, psychology, economics, and power. Moral claims are expressed in humanist or procedural language rather than theological language. Appeals to God, prophecy, discernment, sin, or Logos are treated as private sentiment at best or epistemic contamination at worst.

Friction appears when a figure uses theological grammar as an interpretive tool rather than as mere decoration. References to Logos, spiritual discernment, sin, idolatry, or divine telos can trigger dismissal even if the argument is otherwise careful. Friction also appears when a figure critiques the secular frame as incomplete, or when they treat religious traditions as legitimate sources of moral knowledge rather than as sociological artifacts.

Coding rule in plain terms: if the figure's public reasoning stays within a material public frame and treats theology as non-authoritative, the bit tends toward alignment. If the figure repeatedly uses theological framing as real reasoning and accepts the backlash that comes with it, the bit tends toward friction.

3.5 Church traditionalists and institutional loyalists

This bloc is defined by a binary demand: deference to institutional authority versus prophetic critique of leadership and decline. The key mechanism is boundary policing inside the church: who has permission to criticize, what counts as scandal, and whether unity is defined as truthfulness or as institutional protection.

The alignment signal is rhetorical protection of leaders and structures, preference for internal handling, suspicion of public critique as divisive, and emphasis on submission, order, tradition, and continuity. Even where critique exists, it is framed as minor correction rather than as systemic problem. The institution is treated as the vessel of fidelity, so damaging the vessel is treated as damaging the faith.

Friction appears when a figure publicly names institutional failures, spiritual decline, clerical corruption, or doctrinal drift, especially if they do so in a way that suggests systemic patterns rather than isolated flaws. Friction also appears when a figure implies that institutional loyalty can become idolatry, or when they argue that truthfulness is a higher unity than reputational control.

Coding rule in plain terms: if a figure speaks in ways that protect institutions and treat public critique as suspect, the bit tends toward alignment. If a figure repeatedly triggers defensive reactions for naming structural problems, the bit tends toward friction.

3.6 Charismatic and Pentecostal communities focused on tongues practice

This bloc is defined by a binary demand: affirmation of tongues as a normative sign versus critique of uninterpreted practice and excess. The issue here is not whether charismatic gifts exist. The coding axis is whether the figure's language affirms a particular experiential marker as identity-defining.

The alignment signal includes speaking of tongues as a desirable, expected, or evidential feature of spiritual vitality; treating charismatic manifestations as a major index of God's presence; and interpreting skeptical questions as quenching the Spirit or as intellectual pride. The implicit demand is that experiential display is not merely permitted but prized.

Friction appears when a figure critiques uninterpreted tongues in public worship, emphasizes intelligibility and order, questions performative excess, or warns that certain manifestations can become social signaling rather than spiritual reality. Friction also appears when a figure prioritizes ethical fruit, peace, and disciplined discernment over ecstatic display.

Coding rule in plain terms: if the figure's language normalizes tongues as a boundary marker of authenticity, the bit tends toward alignment. If the figure's language repeatedly sounds like a call for restraint, interpretation, and skepticism about excess, the bit tends toward friction.

3.7 Conspiracy-curious communities and anti-conspiracy skeptics

This axis is intentionally paired because it captures a mirror dynamic. The binary demand is everything-is-a-psyop certainty versus disciplined uncertainty and evidentiary restraint.

Conspiracy-curious audiences often reward pattern detection, hidden-hand explanations, and interpretive certainty under ambiguity. The alignment signal is a readiness to attribute events to coordinated manipulation, to treat official narratives as presumptively false, and to elevate anomaly into proof. They often read "just wait" or "connect the dots" as intellectual courage.

Anti-conspiracy skeptics often reward the opposite certainty: rapid dismissal of non-mainstream claims, preference for institutional sources, and suspicion of pattern narratives as epistemically irresponsible. The alignment signal is a strong emphasis on debunking, procedural verification, and warnings about misinformation contagion.

Friction with both can happen in the same person if they take a third posture: disciplined uncertainty. That posture treats manipulation as possible, sometimes likely, but not as a universal solvent. It insists on evidence, allows "unknown" as an answer, and refuses to turn suspicion into certainty without proof. Conspiracy-curious audiences can call this naive. Anti-conspiracy skeptics can call it irresponsible for even entertaining the possibility.

Coding rule in plain terms: if a figure's language repeatedly leans into high-certainty hidden-hand explanations, conspiracy-curious alignment goes up and skeptic friction goes up. If a figure's language repeatedly leans into debunking certainty, skeptic alignment goes up and conspiracy-curious friction goes up. If a figure consistently holds "manipulation may exist" while demanding proof and tolerating uncertainty, friction can rise on both sides, producing a distinctive pattern.

3.8 AI accelerationists and AI doom or safety absolutists

This axis is also paired because it is structured by opposing binary demands: rapid deployment is salvation versus strong brakes are moral necessity. The model codes where a public figure is experienced as legitimating one of these maximal frames.

Accelerationist audiences reward optimism about deployment, competitive urgency, and claims that benefits will outpace harms. The alignment signal includes dismissing regulation as stagnation, treating adoption as inevitable, emphasizing productivity and growth, and framing caution as fear or as an obstacle to progress.

Doom or safety absolutist audiences reward strong precaution, repeated emphasis on catastrophic risk, and moral urgency for regulation or even moratoria. The alignment signal includes arguing that deployment is reckless, that incentives are misaligned, and that restraint is an ethical demand even at significant cost.

Friction appears with accelerationists when a figure emphasizes limits, misuse, fragility, and the need for constraints, audits, and deliberate governance. Friction appears with safety absolutists when a figure refuses totalizing catastrophe narratives, emphasizes practical mitigation rather than blanket brakes, or treats AI as a tool that can be used responsibly under constraints.

Coding rule in plain terms: if the figure's language reliably sounds like green light and inevitability, accelerationist alignment increases. If it reliably sounds like red light and moral emergency, safety absolutist alignment increases. If it sounds like structured constraint, human accountability, and measured deployment, the figure can generate friction with both maximal camps.

3.9 Academic gatekeepers and anti-academic populists

This axis captures a cultural struggle over epistemic authority. The binary demand is strict disciplinary policing versus public-intellectual insurgency and anti-credential rhetoric.

Academic gatekeepers reward deference to peer review, disciplinary boundaries, technical vocabulary, and credentialed authority. The alignment signal includes "stay in your lane" norms, critique of speculation, and suspicion of interdisciplinary synthesis that does not obey established methods. They often interpret public intellectualizing on technical topics as reckless simplification.

Anti-academic populists reward suspicion of institutions, contempt for credentials, and the claim that "common sense" exposes elite fraud. The alignment signal includes rhetoric about captured universities, ideological gatekeeping, and the dismissal of experts as self-serving.

Friction with gatekeepers appears when a figure speaks across disciplines, rejects certain academic norms, or treats institutions as politically captured. Friction with populists appears when a figure defends expertise, values technical standards, and refuses to treat institutions as purely corrupt. A third posture, again, generates friction with both: respecting method while criticizing institutional capture, insisting on rigor while refusing credential worship.

Coding rule in plain terms: if a figure reliably performs disciplinary deference and treats speculation as suspect, gatekeeper alignment increases. If a figure reliably attacks academia as illegitimate and elevates anti-credential rhetoric, populist alignment increases. If a figure critiques both credential idolatry and anti-intellectual contempt while holding to method, friction can emerge on both sides.

3.10 Nationalist audiences across multiple countries

This axis is defined by a binary demand: my-nation moral exceptionalism versus universal moral language that critiques all sides. Nationalist audiences are not limited to any one country. The relevant feature is the expectation that the speaker will narrate the world with their nation as the moral center and interpret external critique as disloyalty.

The alignment signal is patriotic moral framing that grants the in-group state a presumption of righteousness, interprets geopolitical conflicts as contests between good and evil with clear national heroes, and treats international constraints as either naive or hostile. It often demands selective empathy: full empathy for one's own side, procedural empathy for others.

Friction appears when a figure uses moral language that applies to all states, including their own; treats war crimes, coercion, propaganda, and dehumanization as universal risks; and refuses to turn national interest into moral permission. Friction also appears when a figure's rhetoric foregrounds human dignity across borders, which nationalist audiences can interpret as betrayal of the in-group.

Coding rule in plain terms: if the figure's language reliably centers national righteousness and defends the in-group against moral critique, nationalist alignment increases. If the figure repeatedly applies the same moral constraints to all nations and absorbs "disloyalty" accusations, nationalist friction increases.

In sum, the ten bits are not ten opinions. They are ten recurring pressure points where audiences demand a binary performance. The fingerprint is the pattern of where a public figure reliably affirms those demands and where they reliably irritate them. That pattern can become distinctive in public life, not because it captures the whole person, but because public discourse increasingly forces persons to be legible as codes.

4. Why Ten Bits Can Feel Identifying

At first glance, a 10-bit code seems too small to identify anyone. There are only 1024 possible signatures. In a world with millions of public-facing voices, collisions look inevitable. Yet in practice, people often experience a ten-bit stance profile as "this is obviously that person." The sense of identifiability arises not because 1024 magically becomes large, but because public discourse rarely uses the full 1024 patterns. Social incentives collapse the space. At the same time, certain populations become unusually easy to code. When both processes occur together, a short bitstring can feel surprisingly discriminative.

Three effects explain this: missingness, repeat exposure, and salience. Each effect reduces uncertainty about the code and increases the likelihood that the code functions like a recognizable signature inside a specific reference set.

4.1 The missingness effect

Most people are not fully “bit-coded” because they do not generate public material across all ten axes. They may have no settled view, no interest, or no need to declare anything. Even when they do have a view, they often keep it private or express it only in small, low-visibility contexts that do not trigger wide audience feedback. In terms of data, their stance vector contains many blanks.

This matters because identifiability depends on comparability. A code only distinguishes individuals when many individuals are coded in the same coordinate system. In a general population, most people have missing entries on most axes. If you tried to assign a 10-bit signature to ordinary individuals, you would quickly discover that the more honest representation is not binary but partial: unknown, unstated, inconsistent, or context-specific. The effective code length drops because there is too little observable data.

Public figures occupy the opposite regime. They are repeatedly asked about controversial domains, whether or not they want to be. Their content is recorded, clipped, redistributed, and archived. Their omissions are interpreted. As a result, they are more fully coded, not necessarily because they hold stronger beliefs, but because their role generates a public trace.

This creates a selection effect. The set of people for whom a 10-bit fingerprint feels identifying is not “everyone.” It is a narrower set: those who (a) have a platform, (b) comment across multiple contested domains, and (c) are regularly interpreted by organized audiences. Within that set, missingness shrinks. And once missingness shrinks, the same short code begins to separate people more effectively.

There is a second layer to missingness. Some people have an opinion but avoid stating it to preserve coalition flexibility. That silence itself can be stable and legible. In public discourse, repeated silence on a demanded topic can become a de facto bit: the audience encodes it as fear, complicity, or strategic ambiguity. In other words, missingness does not remain missing; it often gets interpreted into the code. That interpretive filling-in can be unfair, but it is common, and it further reduces the number of “unknowns” for public figures.

So the missingness effect is simple: when most people are un-coded, a code is not discriminative. When a subpopulation is highly coded, the same number of bits becomes more discriminative. Public figures are that highly coded subpopulation.

4.2 The repeat-exposure effect

A single statement is noisy. It can be misunderstood, decontextualized, or unrepresentative. Many stance conflicts are born from first impressions, and first impressions are weak data. Repeat exposure changes that. When a public figure produces a steady stream of content,

ambiguity decreases. The audience sees not just what is said, but what is said repeatedly, what is avoided, what is corrected, what is doubled down on, and what is apologized for.

Repeat exposure turns a stance fingerprint from a speculative inference into a pattern recognition task. The figure's rhetorical habits become visible. Do they insist on symmetrical standards or do they excuse their side? Do they speak in theological frames or in strictly material frames? Do they treat institutions as sacred vessels or as fallible structures open to public critique? Do they treat AI as inevitable salvation or as a reason for strong brakes? Over time, these are not isolated claims. They become a voice.

This effect is amplified by the way modern media recirculates content. Public figures are not merely repeated; they are replayed. Specific phrases become canonical. Opponents and supporters curate compilations that stabilize the figure's "brand." Even if the figure evolves, the archive creates inertia. A person can change their mind, but the platform environment keeps earlier versions of their code alive.

Repeat exposure also forces decision points. Many of the ten axes become unavoidable during crises: a war, an election, a scandal, a technology shock, a cultural flashpoint. When pressure spikes, the figure's response is rarely left ambiguous. It is interpreted, quoted, and used as evidence of alignment or betrayal. Each crisis functions like a calibration event: audiences re-measure the figure's bits.

The result is a stabilizing loop. Public figures speak often. Audiences interpret and react. The figure responds to the reaction, either by clarifying, softening, or hardening. That response creates more data. After enough iterations, the figure's stance fingerprint becomes predictable. People do not merely know what the figure said; they know what the figure will say.

That predictability is what makes ten bits feel like identity. Not because it captures interiority, but because repetition collapses uncertainty. In a high-exposure regime, the bitstring becomes a stable public object.

4.3 The salience effect

Not all topics code equally. Some domains generate weak audience response and therefore weak discrimination. Others generate intense reactions that sharpen the coding boundary. The ten axes chosen in this paper are salient in exactly that sense: they are conflict-rich, identity-linked, and heavily policed by organized audiences.

High-conflict topics behave like high-gain sensors. Small differences in language produce large differences in reaction. A single clause can flip perceived alignment. A refusal to use a demanded slogan can trigger backlash. A single moral distinction can be read as betrayal. This does not mean the audience is correct. It means the audience is attentive and reactive, which is what creates sharp coding.

Salience also reduces the cost of inference. On a low-salience topic, an observer needs many data points to infer stance. On a high-salience topic, the audience has already done the work. Organized blocs maintain interpretive heuristics: lists of approved frames, lists of disallowed frames, and rapid labeling mechanisms. The public figure's stance becomes legible because the audience's reaction is legible. The crowd is not just a receiver; it is a classifier.

This can create a distortion, but it also explains the felt identifiability. When the audience reaction is strong and consistent, it acts like a checksum. People do not only remember the figure's words; they remember the predictable pattern of who gets angry and who applauds. That pattern becomes part of the fingerprint.

Finally, salience concentrates attention on a smaller number of public figures. Many people talk about AI, politics, religion, or geopolitics. Fewer people are widely heard when they do so. The salience of the topics draws attention to prominent voices, and prominence generates more coding data. This creates a funnel. The more salient the topic, the more the conversation centers on a smaller cast, and the easier it becomes for observers to treat a short code as identifying within that cast.

Taken together, these three effects explain the phenomenon. Ten bits are not enough to identify everyone. But in a subpopulation where missingness is low, repetition is high, and salience is extreme, a 10-bit stance fingerprint can feel like a name. The paper's next task is to show the countervailing truth: the same mechanisms that make the code feel identifying also produce predictable collisions, because ideological bundling reduces the effective number of real bits in circulation.

5. Why Ten Bits Often Are Not Identifying

The previous section explained why a 10-bit fingerprint can feel like a name in public discourse. This section explains why that feeling is frequently misleading. The same ecosystem that forces people into bits also compresses those bits into standard bundles. When that happens, the 1024 theoretical signatures collapse into a much smaller menu of socially viable combinations. In other words, the code-space exists on paper, but the culture doesn't actually inhabit it. The result is that many public figures share the same apparent fingerprint, not because they are the same person, but because they are produced by the same coalition logic.

Three mechanisms drive this non-uniqueness: correlation between bits, coalition templates that mass-produce stance bundles, and collisions where distinct people map to the same pattern even when their reasons and lives differ radically.

5.1 Correlated bits and ideological bundling

A ten-bit system only yields 1024 patterns if the bits are genuinely independent. Independence is the exception in politics and culture. In real discourse, many bits are correlated because coalitions bundle positions into packages that signal membership. Once bundling dominates, ten nominal bits can behave like three or four effective bits.

Correlation arises in at least three ways.

First, there is party-line coupling. The partisan bit often drags other bits with it because parties and media ecosystems adopt standard stances on religion in public life, conspiracy discourse, AI governance rhetoric, academic authority, and nationalist framing. A speaker who wants to remain inside a coalition learns, implicitly or explicitly, which combinations "go together." They may not believe the coupling is logically necessary. They comply because it is socially necessary.

Second, there is identity coupling. Some axes are treated as identity markers rather than as policy positions. When that happens, the bit is not chosen issue by issue but inherited with the identity. If you are strongly inside a given cultural identity, certain rhetorical commitments become expected. If you violate them, you risk expulsion. Expulsion risk forces correlation.

Third, there is media coupling. Platforms and formats reward coherent brands. A coherent brand is not the same as a coherent philosophy. It is a consistent pattern that an audience can recognize quickly. Consistency becomes a commodity. That commodity encourages predictable bundling: the "type" must be legible, and legibility is achieved by repeating a known bundle of stances. A person who holds an unusual combination is harder to market, harder to defend, and easier to attack as inconsistent.

The practical implication is severe. If five of the ten bits are tightly correlated, you do not actually have ten degrees of freedom. You have fewer. The code-space collapses. Many signatures become rare or socially impossible. The observed distribution becomes lumpy: a small number of common bundles and a long tail of unusual patterns.

This is why identifiability can fail even in highly observed public-figure sets. The fingerprint might not distinguish a person; it might only distinguish a coalition bundle. What looks like personal identity is often coalition geometry.

5.2 Coalition templates and mass-produced personas

Even when a public figure is sincere, their public persona is not created in a vacuum. It is produced under pressures from donors, party apparatus, activist networks, media training, platform incentives, and audience analytics. These pressures generate templates: standardized rhetorical styles and stance bundles that reliably attract a given audience segment.

A coalition template is not necessarily a conspiracy. It is often emergent. It comes from repetition and selection. Certain packages perform well, so they are promoted. The people who can perform them are invited onto bigger stages. Over time, the persona becomes mass-produced. You can hear the same cadences, the same slogans, the same moral grammar, and the same enemy images across many different individuals.

Templates also arise from professional specialization. Think tanks, advocacy organizations, religious networks, and influencer pipelines train people into recognizable frames. They supply talking points, narratives, and "safe language" that avoids coalition penalties. When a public figure must speak quickly on complex issues, they reach for the template because it reduces risk. It is easier to stay inside the railings than to invent a new track.

This matters for a 10-bit fingerprint because templates generate predictable bit combinations. Once the template is adopted, the bitstring is not a personal signature; it is a uniform. The public figure becomes interchangeable within a category. The "10-bit personality" collapses into a genre.

There is also an adversarial dynamic. Opponents often treat the other side as a template and refuse to perceive individuality. That perceptual refusal can itself cause collisions. Two people with different reasoning can be coded as identical because an out-group audience uses a coarse classifier: "they're one of those." In practice, the fingerprint is sometimes not measuring the figure at all; it is measuring the audience's willingness to see detail.

So coalition templates produce non-uniqueness in two directions: by training many figures into the same packages and by training many audiences to perceive only packages.

5.3 Collisions and look-alikes

A collision occurs when two distinct people map to the same 10-bit pattern. In a strict mathematical sense, collisions are inevitable whenever the population exceeds 1024. But the more important point is conceptual: collisions can occur even in smaller sets because the fingerprint measures friction patterns, not inner logic.

There are several common forms of collision.

First, there are moral collisions. Two people can share the same outward stances for opposite reasons. One person may critique institutional loyalty because they are committed to truthfulness and reform. Another may critique institutional loyalty because they want to burn institutions down. The bitstring will not distinguish them if it only records friction with institutional loyalists. To the audience, both appear as critics. Their motives diverge, but the friction pattern collides.

Second, there are rhetorical collisions. Two people may share a style that triggers similar audience reactions even if their substantive views differ. A sharp tone, an ironic register, or a refusal to use demanded slogans can generate the same backlash from maximalist audiences. The fingerprint then captures style-induced friction, not just content-induced friction.

Third, there are timing collisions. Public figures often shift positions across time. If the code is computed over a window, two different people might match because one is in transition while the other is stable. Alternatively, a single person might collide with their earlier self. The fingerprint is a snapshot, not an essence.

Fourth, there are visibility collisions. A figure who speaks constantly across all axes will be fully coded. Another figure who speaks sparsely may be coded on fewer axes and then "completed" by audience inference. The second figure can be forced into a collision by the audience's assumptions. In that case, the collision is partly an artifact of missingness being filled in.

Finally, there are category collisions. Some audiences react to category membership more than to content. If someone is coded as "religious," "academic," "nationalist," "populist," or "technologist," audiences may pre-assign likely stances and react accordingly. The fingerprint then becomes a reflection of stereotypes rather than of the person's actual speech. Two people who share a category can end up sharing a friction pattern even if they disagree on specifics.

These collisions do not invalidate the model. They clarify what it is good for. A 10-bit audience-fingerprint can be a useful map of how a public figure is positioned in a landscape of coalition pressures. It is less reliable as an identifier of individual identity, and it is dangerously unreliable as a proxy for motives, virtues, or moral worth.

The practical lesson of this section is that the fingerprint must be treated as low-dimensional compression. Compression is valuable because it makes patterns visible, but it always discards information. When observers forget what is discarded, they mistake a type for a person. The model's usefulness depends on keeping that remainder in view.

6. Coding Method: Turning Public Output into a Bitstring

If "the 10-bit personality" is an audience-fingerprint, then the central technical question is simple: how do we code it without turning it into a Rorschach test? The danger of stance-fingerprinting is not only moral misuse. It is methodological sloppiness. If the coding rules are vague, every observer will "find" the bitstring they expected to find. The model then becomes a mirror of bias rather than a map of public pressures.

This section proposes a practical coding method built around four principles. First, code public output, not imagined interior beliefs. Second, privilege long-form over sound bites. Third, use

thresholds and repetition to avoid single-quote overfitting. Fourth, explicitly represent uncertainty by attaching a confidence score to each bit. The goal is not to produce a final verdict; it is to produce a disciplined, revisable stance profile.

6.1 Data selection rules

The unit of evidence is public output that can reasonably represent the figure's considered stance or recurring rhetorical posture. Not all public speech is equal. A good coding set prioritizes sources that are (a) attributable, (b) context-rich, and (c) stable enough to survive reasonable interpretation.

What counts as evidence in this method:

- Speeches and prepared remarks that are publicly available in full form, ideally with transcript-level text.
- Essays, op-eds, newsletters, blog posts, and published long-form writing where the figure controls the wording.
- Long-form interviews (podcasts, print interviews, extended video) where the figure is asked follow-up questions and has space to qualify.
- Public debates with sufficiently long answers to show the figure's actual reasoning rather than a single slogan.
- Official statements: press releases, organizational statements signed by the figure, or public letters.

What is weak evidence and should be de-weighted:

- Short clips, viral excerpts, or meme-captions divorced from the full context.
- Anonymous summaries of what the figure "said," especially when no primary text is available.
- Second-hand accounts, even if widely repeated.
- Social media one-liners, unless they are repeated and consistent enough to represent a stable posture.
- Questions answered under obvious coercion or extreme time pressure, where the answer functions as crisis management rather than stance.

A practical rule is to assemble a "coding corpus" of at least three distinct long-form items per axis whenever possible. If the figure does not speak in long form, the method should admit that the bit is hard to code. Silence is data, but it is not the same as a stance.

A second practical rule is adversarial sampling: include at least one item where the figure is speaking to an out-group or being challenged by an interlocutor who does not share their assumptions. In-group speech can be performative; cross-pressure speech reveals the shape of what a figure will defend.

6.2 Coding rules per bit

Each bit in this model is a friction-versus-alignment axis with a specific audience bloc. That suggests a straightforward coding question: does the figure reliably produce language that this bloc experiences as aligned and affirming, or as irritating and disloyal? But "reliably" is the keyword. The coding must guard against three failure modes: overfitting to a single quote, confusing style for stance, and confusing audience misreading for the figure's posture.

A minimal set of coding rules can keep the method honest.

Rule 1: Use repetition thresholds

A bit should not be set based on a single statement unless the statement is a formal declaration that clearly commits the figure. Otherwise, require repetition: at least two to three separate instances across different contexts where the same stance posture appears. If the figure is consistent only in one setting and inconsistent elsewhere, the bit should remain ambiguous or low-confidence.

Rule 2: Distinguish content from meta-stance

Some bits are triggered by content, others by the figure's refusal to accept the binary demand itself. For example, friction with hard partisans often shows up as a meta-stance: rejecting team logic, insisting on symmetrical standards, or refusing purity tests. In those cases, code the meta-stance rather than trying to infer hidden policy detail.

Rule 3: Avoid single-quote overfitting

Overfitting happens when one vivid line becomes the whole identity. To prevent it, require that the line is representative of a broader pattern. Ask: does the figure elsewhere double down, clarify in the same direction, or incorporate the same framing repeatedly? If not, treat it as noise.

Rule 4: Handle ambiguity explicitly

Ambiguity is not a failure. It is a legitimate state of the data. If a figure uses balanced language that could be read as aligned or disloyal by different sub-factions, code the bit as indeterminate or mixed rather than forcing a binary decision. The goal is not a perfect bitstring; the goal is a faithful one.

Rule 5: Use audience reaction as supporting, not primary, evidence

Audience reaction can be informative, but it is not decisive. Some audiences misread. Some audiences choose to attack for strategic reasons. Reaction is best used as corroboration: if the

figure repeatedly triggers predictable backlash from a bloc while also expressing the content that would reasonably cause it, confidence rises. If the backlash occurs without corresponding content, treat it cautiously.

Rule 6: Decide the direction of the bit in advance
Each bit needs an explicit convention, such as: 1 means "frequently produces friction with this bloc," 0 means "frequently produces alignment/affirmation with this bloc." The convention must be fixed before coding to prevent post hoc reinterpretation.

Rule 7: Separate mixed cases from binary cases
Many figures will be mixed: they align with the bloc on one subtopic but irritate it on another. In that case, the method should either (a) keep the bit but lower confidence, or (b) break the axis into sub-bits in a larger model. For this paper's 10-bit discipline, the right move is usually to keep the bit but mark it as mixed with a low confidence score.

A useful practical tool is a short "justification note" per bit: two to four sentences that cite the recurring pattern (not citations, just internal reasoning), and name what would flip the bit if new evidence appeared. This prevents the code from becoming a frozen accusation.

6.3 Confidence scoring

Binary coding is attractive because it is simple. The problem is that public life is noisy, strategic, and full of missingness. A purely binary output pretends to certainty that often does not exist. Confidence scoring solves this by pairing each bit with a companion value that tracks the coder's certainty.

A simple confidence scale that remains readable:

- 0: Uncoded (insufficient data)
- 1: Low confidence (few data points, high ambiguity, strong context dependence)
- 2: Medium confidence (multiple data points, some ambiguity remains)
- 3: High confidence (repeated across contexts, stable framing over time)

This companion score accomplishes several things.

First, it prevents false precision. A figure can have a bit set while still being coded as low confidence, which communicates that the fingerprint is provisional. Second, it allows meaningful comparison between figures who are highly visible and figures who are selectively visible. Third, it makes drift visible. A figure whose bit flips with low confidence might be noise. A figure whose bit flips after a period of high confidence signals a real change.

Confidence can be raised by: repetition across time, consistency across in-group and out-group contexts, explicit endorsements or explicit rejections of the bloc's demanded frame, and stable response under crisis pressure. Confidence can be lowered by: one-off statements, forced clarifications that retreat from the frame, evidence of strategic ambiguity, and high variance across contexts.

The key discipline is to treat confidence as part of the model, not as an afterthought. A bitstring without confidence is a caricature engine. A bitstring with confidence begins to behave like an epistemically responsible map.

6.4 Time and drift

Public stance is not static. People change. Context changes. Coalitions change. Certain audience blocs become more or less demanding. A stance fingerprint is therefore a time-indexed object. The same person can carry different signatures across different years, and the same signature can mean different things depending on the surrounding environment.

Two ideas are crucial: drift and recency windows.

Drift is the gradual change of a signature across time. It can be caused by genuine intellectual evolution, spiritual conversion, new institutional roles, strategic repositioning, or reaction to audience punishment and reward. Drift can also occur when the environment changes: a war begins, a platform changes its incentives, a movement radicalizes, or a formerly fringe view becomes mainstream. In those cases, a person can keep saying the same thing while the audience flips from affirmation to hostility. The bit flips even if the speaker stays constant, because the bit codes friction patterns.

Recency windows matter because public reputations are path-dependent. A signature computed over a decade may average out meaningful changes. A signature computed over a month may overreact to a crisis. A practical approach is to compute fingerprints over multiple windows: a long window to capture baseline stance, and a short window to capture current posture. The difference between them is itself informative: it shows whether the figure is stable, drifting, or in a phase transition.

A disciplined method can state a default window, such as the last two to three years for "current signature," and the last eight to ten years for "career signature," depending on the figure's activity. For public figures who have recently surged into prominence, the window should be shorter but still large enough to avoid overfitting to a single event.

Finally, time discipline also protects ethics. Many reputational harms come from freezing a person at their worst moment or from treating an old stance as permanent identity. Drift-aware coding reminds the reader that the fingerprint is a snapshot, not an essence.

In summary, turning public output into a bitstring is less like reading someone's mind and more like mapping a dynamic interaction between a speaker and a pressure-filled audience ecosystem. The coding method must therefore be explicit, repetition-based, confidence-scored, and time-indexed. Without those constraints, a 10-bit fingerprint becomes a weapon. With them, it can remain a tool for clarity about how public discourse compresses people into legible codes.

7. Worked Examples Without Name-Claims

A 10-bit audience-fingerprint becomes persuasive only when it feels usable. But "usable" is exactly where misuse begins: once a model looks like a classification tool, people instinctively apply it to real individuals and then treat the resulting bitstring as a verdict. This paper deliberately refuses that move. The aim is to demonstrate discrimination power and collision risk without providing a doxxing-by-inference mechanism or inviting readers to turn the fingerprint into an accusation machine.

So this section uses archetypes: stylized, recognizable patterns that appear across many public figures. Each archetype is presented as (a) a narrative description of the public posture, (b) an illustrative 10-bit pattern, and (c) a note on what tends to raise or lower confidence per bit. The patterns are not meant to be exhaustive, and they are not moral rankings. They are maps of how different public personas interact with the ten audience blocs.

For simplicity, we use the convention introduced earlier: for each axis, 1 means the figure tends to generate sustained friction with that audience bloc, and 0 means the figure tends to generate alignment or affirmation from that bloc. In practice, each bit should be paired with a confidence score, but for readability we describe confidence in prose.

The ten axes appear in the same order throughout:

1. U.S. hard partisans, 2) pro-Israel maximalists, 3) pro-Palestinian maximalists, 4) strong secularists, 5) church institutional loyalists, 6) charismatic tongues-focused communities, 7) conspiracy-curious plus anti-conspiracy skeptics (as a paired axis where friction can arise with one side or the other depending on posture), 8) AI accelerationists plus AI doom or safety absolutists (paired), 9) academic gatekeepers plus anti-academic populists (paired), 10) nationalist audiences.

These archetypes are intentionally broad. Their value is that many real people can resemble them without being reducible to them.

7.1 The partisan amplifier

This archetype is defined by reliable team reinforcement. Their role is not merely to hold partisan opinions but to amplify partisan identity through language that flatters the in-group and degrades the out-group. They tend to avoid cross-tribal symmetry, avoid acknowledging weaknesses on their own side, and treat moral evaluation as a tool of coalition warfare rather than as a universal standard.

In public life, the partisan amplifier often appears as a highly consistent brand. They rarely surprise their base. They speak in a way that makes their allies feel safe and makes their opponents feel caricatured. When faced with an inconvenient fact, they pivot to a larger narrative that preserves team innocence.

Illustrative fingerprint (conceptual):

- 1. Partisans: 0 (alignment)
- 2. Pro-Israel maximalists: depends on coalition; often 0 (alignment) or stable 1 (friction) based on party-media ecosystem
- 3. Pro-Palestinian maximalists: often the complement of bit 2 in their discourse environment
- 4. Strong secularists: depends; many partisan amplifiers mirror their coalition's religious style rather than a principled stance
- 5. Church institutional loyalists: if their base includes strong church institutionalism, tends toward 0
- 6. Charismatic tongues communities: often not directly engaged; if engaged, usually matches coalition religious style

7. Conspiracy-curious vs anti-conspiracy skeptics: tends to align with whichever side benefits the coalition narrative at the time

-

8. AI maximal commitments: often opportunistic; aligns with coalition mood rather than building a stable AI ethic

-

9. Academic gatekeepers vs anti-academic populists: frequently 1 with gatekeepers and 0 with populists in populist coalitions; reversed in technocratic coalitions

-

10. Nationalists: often 0 (alignment), because nationalist moral framing pairs well with partisan identity

The key property is not the exact bitstring but low friction with coalition enforcers. This archetype's "identity" is that they rarely collide with the audiences that sustain them.

What raises confidence: repeated use of loyalty language, predictable enemy images, consistent selective outrage, punishment of in-group dissent.
What lowers confidence: episodes of genuine cross-tribal critique that persist under backlash, or stable universal moral standards that override team advantage.

Collision profile: very high. This archetype produces many look-alikes because it is template-driven. Many individuals can share this fingerprint while differing in biography and temperament.

7.2 The prophetic institutional critic
This archetype is defined by conflict with institutional protection logic, especially in religious contexts. The prophetic critic speaks as if truth-telling is a moral duty even when it embarrasses leaders. They often name spiritual decline, leadership failure, moral compromise, or institutional idolatry. They may still love the institution, but they refuse to protect its reputation at the expense of integrity.

Unlike the partisan amplifier, this archetype often generates friction from inside their own camp. Their critique is experienced as disloyal by institutional loyalists, even when the critic claims loyalty to the faith rather than to the administration.

Illustrative fingerprint (conceptual):

- 1. Partisans: can be mixed; prophetic critics sometimes resist team logic, sometimes use it instrumentally
- 2. Pro-Israel maximalists: mixed; depends on whether their moral framework constrains state power
- 3. Pro-Palestinian maximalists: mixed; they often resist totalizing condemnation and dehumanization
- 4. Strong secularists: often 1 (friction), because prophetic language is unapologetically theological
- 5. Church institutional loyalists: 1 (high friction)
- 6. Charismatic tongues communities: mixed; they may affirm gifts but critique performative excess and uninterpreted public display
- 7. Conspiracy-curious vs anti-conspiracy skeptics: often friction with both extremes, depending on their discipline
- 8. AI maximal commitments: often neither camp; they may treat AI as a moral test rather than as salvation or doom
- 9. Academic gatekeepers vs anti-academic populists: mixed; prophetic critics may distrust institutional authority broadly but still respect method

10. Nationalists: often 1 (friction) if they apply universal moral critique to their own nation

The defining bit is 5: sustained friction with institutional loyalists. Everything else can vary. That makes this archetype discriminative on certain axes but also prone to misclassification if observers confuse “public critique” with “lack of faith.”

What raises confidence: long-form critique, explicit naming of institutional failure, acceptance of backlash without retreat, consistent appeal to higher standards than reputation.
What lowers confidence: occasional scandal commentary without sustained pattern, or criticism that is purely partisan rather than morally grounded.

Collision profile: moderate. Many different theological traditions produce their own prophetic critics, so collisions are common across denominations. Motives can also diverge: reformers and wreckers can look similar in a 10-bit compression.

7.3 The technologist with AI maximal commitments
This archetype is defined by a strong, consistent commitment to an AI maximal frame, either accelerationist or safety-absolutist. The key is not “interest in AI,” but an identity-laden posture that treats AI as the central civilizational axis and demands moral urgency in one direction.

There are two common subtypes.

Subtype A: acceleration maximalist. They interpret AI as inevitable progress and treat restraint as cowardice or sabotage. They emphasize competition, speed, and opportunity, and they discount harms as either manageable or outweighed.

Subtype B: safety maximalist. They interpret AI as a catastrophic-risk engine and treat deployment as moral negligence. They emphasize incentive misalignment, irreversible harm, and the need for strong brakes.

In both subtypes, the AI bit becomes disproportionately informative because it drives alliances and conflicts with other blocs.

Illustrative fingerprint tendencies:

- 8. AI maximal commitments: 0 with accelerationists / 1 with safety absolutists, or the reverse depending on how the paired axis is coded in practice
-

9. Academic gatekeepers: safety maximalists often align more with academic or scientific authority; acceleration maximalists often align more with entrepreneurial anti-gatekeeping rhetoric

- 10. Nationalists: acceleration maximalists may adopt nationalist competition frames; safety maximalists may adopt global-risk frames that irritate nationalists

- 7. Conspiracy vs skeptic: both subtypes can be misread as conspiratorial if they emphasize hidden risks or hidden capabilities, or misread as naive if they underplay them

This archetype shows why collisions are predictable. Many technologists share similar incentives, talk tracks, and coalition alignments. The fingerprint frequently identifies a genre rather than a person.

What raises confidence: repeated, explicit AI-first framing across years, consistent policy stance, refusal to moderate under backlash.
What lowers confidence: opportunistic AI rhetoric that tracks news cycles rather than a stable worldview.

Collision profile: high. In a compressed stance space, AI maximalists tend to cluster into a few recognizable bundles. Their fingerprints collide often even when technical sophistication varies.

7.4 The cross-pressure peacemaker
This archetype is the one most likely to produce a distinctive pattern within public-figure sets, not because it is morally superior but because it refuses multiple maximal demands at once. The cross-pressure peacemaker tends to speak in universal moral language and insists that compassion and constraint apply to all sides. They critique institutional idolatry, resist partisan team logic, refuse dehumanizing rhetoric, and insist on evidentiary discipline in a rumor economy. They may be religious or not, but they resist frames that collapse the moral universe into a single tribe's victory story.

Because this archetype disappoints many gatekeepers at once, it can produce high friction across multiple blocs. That friction pattern can become discriminative because it is less common: most public figures choose a stable patron audience and optimize for that audience's approval. The peacemaker does not optimize that way.

Illustrative fingerprint (conceptual):

- 1. Partisans: 1 (friction)

-

2. Pro-Israel maximalists: 1 (friction) if the peacemaker insists on constraints and names Palestinian suffering as morally weighty

-

3. Pro-Palestinian maximalists: 1 (friction) if the peacemaker refuses totalizing condemnation and refuses dehumanization

-

4. Strong secularists: depends; if theological language is present, friction rises; if not, it may be neutral

-

5. Church institutional loyalists: often 1 (friction) if the peacemaker critiques institutional protectionism

-

6. Charismatic tongues communities: mixed; if the peacemaker emphasizes intelligibility and restraint, friction rises

-

7. Conspiracy-curious and anti-conspiracy skeptics: often friction with both extremes if the posture is disciplined uncertainty

-

8. AI maximal camps: often friction with both extremes if the posture is structured constraint rather than salvation or doom

-

9. Academic gatekeepers and anti-academic populists: often friction with both if the posture is respect for method plus critique of institutional capture

-

10. Nationalists: often 1 (friction) if universal moral critique is applied to the in-group nation

What makes this archetype distinct is that it accumulates friction not by being contrarian, but by maintaining invariants that do not serve any one maximal bloc's identity needs.

What raises confidence: repeated symmetrical moral constraints, careful language hygiene, resistance to demanded slogans, consistent insistence on evidence, and refusal to treat criticism as betrayal.
What lowers confidence: occasional “both sides” language without substantive constraint, or neutrality that avoids moral clarity.

Collision profile: lower than the other archetypes but not zero. Some figures adopt cross-pressure rhetoric as branding while quietly signaling to a base. Others genuinely hold the posture but differ in theology, temperament, or policy specifics that the bitstring cannot capture.

These worked examples show both the promise and the limit of the method. A 10-bit audience-fingerprint can discriminate between common public personas, and in some cases it can feel near-unique. But the same compression that makes it readable also guarantees collisions. The fingerprint is therefore best used as a map of public cross-pressures, not as an identity replacement and not as a moral verdict.

8. Interpretation Limits

A 10-bit audience-fingerprint is a deliberate compression. Its value is that it makes the “shape” of a public figure’s friction landscape visible. Its danger is that it invites readers to treat the compression as the person. The most important intellectual work, therefore, is not generating the bitstring but interpreting it with restraint. The model is only responsible if it comes with clear limits on what can be inferred.

Three interpretation limits matter most. First, positions do not uniquely determine reasons. Second, audience reaction does not uniquely determine author intent. Third, the person always exceeds the code.

8.1 Positions versus reasons

A bitstring records an outward pattern: alignment or friction with an audience bloc. Even if we translate that into “positions,” the model still does not tell us why those positions are held. Two people can share the same fingerprint while standing on opposite moral ground.

This happens because the same outward stance can be reached through incompatible ethical frameworks.

One person might critique institutional loyalty because they believe truthfulness is sacred and reform is necessary for the institution’s survival. Another might critique institutional loyalty because they reject all institutions and prefer permanent disruption. Both produce friction with

institutional loyalists. The bit cannot tell whether the critique is a surgeon's cut or an arsonist's match.

Similarly, a person might resist partisan team logic out of a principled commitment to universal standards. Another might resist team logic because they are opportunistic and want to keep multiple audiences available. Both irritate hard partisans. The bit cannot tell whether the refusal is conscience or strategy.

On Israel–Palestine, identical friction patterns can arise from sharply different theologies and moral grammars. A universalist peacemaking ethic can produce friction with maximalists on both sides. So can a cynical "everyone lies" posture that refuses commitment. The bitstring captures the same alignment pattern while the underlying spiritual posture is almost opposite.

The same is true for secularism versus religious framing. A figure might use theological language because they believe it names reality and grounds moral truth. Another might use it as mere cultural signaling. Both irritate strong secularists. The bit cannot detect sincerity.

The deeper point is that a stance fingerprint has low semantic resolution. It is closer to a silhouette than a portrait. It can show you the outline of public positioning, but it cannot show you the inner architecture: the chain of reasons, the hierarchy of values, the experience that formed the person, or the discipline that restrains them.

This is why the model must never be used as a shortcut for moral judgment. The same bits can arise from love or from contempt, from humility or from vanity, from disciplined courage or from reckless impulse. Without reasons, you do not actually know what you are looking at.

8.2 Audience reaction versus author intent

The model codes friction patterns, and friction is relational. It happens between a speaker and an audience, within a particular context, under particular incentives. That means friction can be generated by substance, but it can also be generated by misunderstanding, misquotation, strategic distortion, and audience projection.

Three dynamics make this especially important.

First, audiences misread. People hear what they expect to hear. If a figure has been pre-labeled as a type, then ambiguous statements get decoded to fit the label. A person who is coded as "religious" can be dismissed as irrational even when they make a careful empirical argument. A person coded as "academic" can be dismissed as elitist even when they speak plainly. These misreadings can generate friction independent of what was actually said.

Second, audiences distort strategically. In high-conflict domains, backlash is not always a spontaneous emotional response; it can be a deliberate tactic. A coalition may attack a figure not because the figure is truly disloyal, but because attacking signals loyalty to the base and

deters others from deviating. In that case, friction is partially manufactured. The fingerprint will still capture it, but the meaning shifts: the bit reflects the coalition's discipline mechanisms as much as the speaker's content.

Third, media formats decontextualize. A figure can speak in long-form nuance and still be represented by a clipped line that reverses their meaning. Once that clip circulates, the audience reacts to the clip, not to the person's considered stance. The speaker's intent is now downstream of the ecosystem's editing. The friction becomes real even if it is based on a false representation.

These dynamics create a practical rule for interpretation: never treat backlash as proof by itself. Backlash can corroborate a pattern, but it cannot replace reading the primary output. The model is safer when it codes the speaker's recurring rhetorical commitments and uses audience reaction as secondary confirmation, not as primary evidence.

More importantly, intent and impact can diverge. A person can intend peace and produce rage. A person can intend provocation and produce admiration. A person can intend honest inquiry and be received as hostile. Because the fingerprint is about reception, it can mislead you about intent unless you also study the person's long-form reasoning, tone, and corrective behavior over time.

8.3 The person exceeds the code

Even if we coded perfectly, the bitstring is still a compression of public stance. A person is not identical to their public stance. The code compresses away everything that is not directly visible in public output.

What it compresses away includes virtues. Courage, patience, humility, truthfulness, and mercy are not straightforwardly readable from a stance vector. A person might hold a correct position without virtue, or hold an incorrect position with a remarkable sincerity and charity that makes reform possible. The fingerprint cannot see that.

It compresses away temperament. Some people are naturally sharp, others gentle. Some are conciliatory, others confrontational. Temperament shapes how words land. Two people can say the same thing and generate different friction because their manner differs. The bitstring will treat them as identical, but the social reality is not identical.

It compresses away private charity and private suffering. Many public figures are far kinder in private than their public persona suggests, or far more broken than their public performance reveals. Some are constrained by responsibilities, health, trauma, or family realities that change what they can say and how they say it. The fingerprint knows none of this.

It compresses away the inner life: prayer, repentance, fear, joy, doubt, spiritual struggle, and the quiet motives that make someone faithful or faithless in the most important sense. If you care about the human being, the inner life matters more than the stance vector. A code can describe a public signature while missing the soul.

This is the final interpretation limit: the bitstring is a tool for mapping public cross-pressure and coalition dynamics, not a substitute for knowing a person. If you use it as a substitute, you will dehumanize. If you remember it is compression, you can gain clarity without losing dignity.

A good way to keep this limit active is to pair every fingerprint with an explicit humility clause in your own mind: “This code describes how audiences collide with this speaker in public. It does not tell me the speaker’s full reasons. It does not tell me their intent. It certainly does not tell me their worth.” If that clause is not present, the model becomes spiritually and intellectually unsafe.

9. Ethical Constraints and Misuse Risks

A stance fingerprint is a powerful object because it is small. Small objects travel. They can be repeated, remembered, and deployed as weapons. A bitstring feels clean and decisive in a messy world, which is exactly why it can become morally dangerous. The same compression that makes the model useful for mapping public cross-pressures also makes it easy to misuse for dehumanization, reputational harm, and automated stereotyping.

This section states the ethical constraints that must accompany any use of a 10-bit audience-fingerprint. The constraints are not decorative. They are part of the model. Without them, the model functions as a technology of simplification that accelerates the very moral failures it claims to diagnose.

9.1 Compression and dehumanization

Compression is never neutral. When we compress a person into a short code, we implicitly change what we think a person is. The temptation is to treat a human being as a stance bundle, and then treat stance bundles as moral essences. That is reductionism. And reductionism invites caricature.

The pathway from compression to dehumanization is predictable.

First, the code becomes a substitute for attention. Instead of reading carefully, listening, and distinguishing, we look at the code and assume we already know. This produces epistemic laziness. But the moral problem is deeper: attention is a form of respect. When attention collapses, respect collapses.

Second, the code becomes a purity test. Once someone is represented as bits, others begin to demand bit-flips as proof of loyalty. The person becomes a set of levers to pull, not a mind to engage. In that regime, growth is punished because it looks like inconsistency. Nuance is punished because it looks like evasion. Repentance is punished because it looks like weakness. Compression hardens identity into a prison.

Third, the code becomes an enemy mask. A compressed representation makes it easier to hate. It removes the "face." Real humans are difficult to hate at close range because they contain contradictions, vulnerability, and ordinary needs. A bitstring contains none of that. It is an abstraction. Abstractions are easy to attack.

Fourth, the code invites essentialism. Observers begin to talk as if the code is what the person is, rather than what the person said under specific conditions. The stance becomes fate. The person becomes a type. And types are disposable.

These are not theoretical risks. They are already present in everyday political and cultural speech. People routinely reduce others to a few signals, then treat those signals as total identity. The 10-bit model, if used carelessly, can become a sophisticated version of an existing vice.

So the ethical constraint here is blunt: never present a stance fingerprint as a full description of a person. Never present it as a measure of moral worth. Never forget that a code is a tool for the observer, not a verdict on the observed.

9.2 Reputation harm and automated stereotyping

A stance fingerprint becomes especially dangerous when it is used as a label in social systems. Labels can be used to exclude, to target, to punish, or to justify violence. Even when the stakes are lower, labels can produce real harms: job loss, social isolation, harassment, and long-term reputational damage.

There are several specific misuse risks.

One risk is doxxing-by-inference. Even if a model avoids naming individuals, a sufficiently distinctive code paired with a narrow population can allow observers to triangulate. People love puzzles. A bitstring can become a "guess who" game. In volatile environments, that game becomes threat.

Another risk is guilt-by-pattern. Observers may treat a fingerprint as proof of hidden affiliations. "If your bits look like theirs, you must be one of them." This is not reasoning. It is associational contamination. It collapses difference and punishes people for being classifiable.

A third risk is automated stereotyping. Once a person can be represented as a vector, systems can act on that vector. Employers, platforms, donors, and institutions can use stance fingerprints

to filter, rank, suppress, or amplify. The system does not need to understand the person; it only needs to sort them. Sorting then becomes a mechanism of control.

A fourth risk is adversarial manipulation. Once people know they are being fingerprinted, they can game the model: perform the expected bits to avoid penalties, hide real beliefs behind ambiguity, or bait opponents into misclassifying them. The fingerprint then becomes part of the propaganda environment rather than a tool for clarity. It can be weaponized to frame someone as disloyal by selectively presenting outputs that flip the appearance of a bit.

A fifth risk is freezing and permanence. Digital archives do not forget. A bitstring attached to a person can outlive their actual beliefs. People mature, repent, learn, and change. A reputation label can persist as if change is impossible. This creates moral despair and discourages honest revision.

Finally, there is the risk of asymmetric harm. Powerful figures can survive caricature; ordinary people often cannot. A stance fingerprint method that is casually applied "downward" becomes an instrument of intimidation. The model can tempt observers to classify private individuals who never consented to being treated as public outputs.

So the ethical constraint is this: stance fingerprints must not be used as targeting tools, employment filters, exclusion mechanisms, or harassment accelerators. If a reader's primary intention is to punish or sort people, the model should be treated as off-limits.

9.3 A humane protocol

A humane protocol is a set of rules that keeps a stance fingerprint from becoming a dehumanization machine. The protocol does not guarantee virtue, but it forces certain forms of humility and care into the process.

Here is a practical protocol for responsible use.

Rule 1: Provisional coding only

Treat every bitstring as a snapshot, not an essence. State the time window. Assume drift is possible. Do not speak as if the code is permanent identity.

Rule 2: Explicit uncertainty

Attach confidence to every bit. If confidence is low, say so. If the figure does not speak on an axis, do not fill the gap with assumption. Do not treat missingness as guilt.

Rule 3: Primary output first

Do not code from reactions alone. Read or listen to the long-form output. Use audience reaction as secondary evidence.

Rule 4: Separate description from evaluation
A fingerprint describes how a figure collides with blocs. It is not a moral scorecard. If you want to evaluate, you must do a separate argument: what is true, what is just, what is loving, what is wise. Do not smuggle evaluation into classification.

Rule 5: Do not publish fingerprints for private individuals
Public-figure methods should remain in the public-figure domain. Using stance fingerprints on private people is a fast path to harassment and coercive control.

Rule 6: Avoid naming targets in volatile contexts
Where there is a foreseeable risk of harassment, treat the method as a general analysis tool, not as a person-identification tool.

Rule 7: Allow reasons to matter
If you are tempted to use a bitstring as a verdict, stop and ask: what are their reasons? Have I read them? Have I represented them fairly? If not, the code is not enough.

Rule 8: Maintain the human remainder
When speaking about someone, include at least one explicit reminder that the person exceeds the code: that they may have virtues you do not see, that they may have constraints you do not know, that they may be acting under pressure.

Rule 9: Refuse to equate a code with a soul
This is the deepest rule. The code is about public friction patterns. A soul is about personhood, responsibility, repentance, love, and the hidden life. Do not confuse legibility with moral truth. Do not confuse classification with judgment.

If these rules feel heavy, that is because they should. Classification is power. Even a small model is a lever. The more a culture is addicted to simplification, the more dangerous simplification tools become. The best use of the 10-bit model is therefore not to accelerate tribal sorting, but to reveal the pressures that force sorting in the first place, and to defend the dignity of persons against the cruelty of compression.

10. Applications

A 10-bit audience-fingerprint can be misused as a labeling weapon, but it can also be used as a clarity tool—especially when the goal is not to identify a person, but to understand the pressure landscape that shapes public speech. The most defensible applications treat the fingerprint as an instrument for diagnosing discourse dynamics: where binary demands are being imposed, where coalitional incentives are driving distortion, and where a speaker is being forced into a role that may not match their deeper reasoning.

This section sketches four application domains that benefit from the model's compression while remaining constrained by the ethical protocol described earlier.

10.1 Journalism and debate moderation

Journalism and debate formats often unintentionally intensify binary outputs. The typical interview structure is already a bit-extraction machine: "Do you condemn X?" "Are you for or against Y?" "Will you say the slogan?" When a moderator rewards clarity and punishes nuance, they train public figures to speak in the very performative binaries that poison public understanding.

Used carefully, the 10-bit model can help journalism do the opposite: clarify coalitional pressures without rewarding the performance of coalition loyalty.

One use is pre-interview preparation. A journalist can use the ten axes to anticipate which audience blocs will interpret the conversation as a loyalty test and which questions are likely to be forced into binary traps. That awareness allows the interviewer to design questions that reveal reasoning rather than extracting slogans. Instead of "Which side are you on?" the interviewer can ask "What moral constraints do you apply to your own side?" Instead of "Will you condemn?" the interviewer can ask "What do you think condemnation accomplishes here, and what concrete standard are you applying?"

Another use is moderation that explicitly resists forced binaries. A debate moderator can name the pressure without endorsing it: "Many audiences will demand a yes/no here. Before we go there, state the principle you're applying." This does not remove the audience pressure, but it changes the incentive. It rewards the articulation of standards and reasoning rather than tribal alignment.

A third use is framing and post-debate analysis. Journalists can describe how audience blocs attempted to enforce bits—through backlash, clipped quotes, and loyalty accusations—without treating that backlash as decisive evidence of the speaker's moral standing. Instead of reporting the backlash as the story, journalism can report the coercion pattern as the story: which groups demanded which performances, and what that demand reveals about the discourse environment.

Finally, the model can help avoid false equivalence. The goal is not to say "both sides are equally extreme." The goal is to see where maximalist demands exist and how they operate. Some contexts are asymmetrical. Some blocs are more disciplined or more punitive. The 10-bit map can help moderators recognize that asymmetry without collapsing into propaganda.

In short, journalism can use the model as a diagnostic tool to redesign conversation formats so they extract principles, not bits—while still acknowledging that audiences will attempt to re-encode everything into bits after the fact.

10.2 Computational social science and discourse mapping

In computational social science, stance vectors are already common. What this paper adds is a specific, audience-pressure framing and a caution about effective information content. Ten nominal bits do not equal ten effective bits when correlation is high. The key analytic tasks are therefore measuring effective bits, measuring correlation structure, and tracking drift over time.

One application is estimating effective dimensionality. In practice, many bits co-move because coalitions bundle positions. A researcher can treat observed bitstrings as data and ask how many independent degrees of freedom they contain. If a population's stance fingerprints cluster tightly into a handful of patterns, the effective dimensionality is low. That tells you something important: the discourse system is not supporting independent thought across axes; it is supporting packaged identity.

Another application is correlation mapping. Researchers can measure which bits are tightly coupled and which remain relatively independent. That coupling map is more informative than the raw fingerprints, because it reveals how coalitional logic binds disparate issues together. For example, if religious framing is highly coupled with nationalism and anti-academic populism in one media ecosystem, that tells you how that ecosystem constructs identity. In a different ecosystem, the couplings might reverse. The model thus becomes a comparative tool for different public spheres.

A third application is drift detection. If you compute fingerprints across rolling time windows, you can track how individuals, movements, or entire populations move. Drift can be personal (a figure changes), contextual (the environment changes), or classificatory (audiences reinterpret stable speech differently). Time-indexed fingerprints can help disentangle these: did the person change their rhetoric, or did the audience's boundary shift?

A fourth application is polarization measurement that is more granular than left/right. Traditional polarization metrics compress politics into a single line. The 10-bit framework produces a multi-axis map that can show whether polarization is "one-dimensional" (everyone follows party bundle) or "multi-dimensional" (new coalitions, cross-cutting commitments, and unusual combinations). That matters for predicting instability. Highly bundled systems are brittle. Cross-cut systems can be noisy but sometimes more resilient.

These computational uses must remain ethically constrained. The point is to study discourse structure, not to build a targeting database. The highest-value outputs are aggregate and structural: correlation networks, clustering patterns, drift curves, and measures of bundling strength—not lists of labeled individuals.

10.3 Personal authorship and self-consistency audits

The most constructive use of the 10-bit model may be personal: using the ten axes as a mirror rather than as a weapon. In that mode, the model becomes an authorship audit tool. It helps a writer see where they are being pulled into performative binaries, where they are inadvertently signaling to blocs they do not wish to serve, and where their stated principles fail to constrain their speech consistently.

A self-audit starts by choosing the time window and assembling one's own corpus: essays, interviews, posts, talks, and any public writing. Then the author codes themselves, ideally with confidence scores, and asks three questions.

First, are my bits consistent with my stated moral invariants? If I claim to apply universal moral constraints, do I actually do it when my in-group is implicated? If I claim to resist maximalism, do I resist it on both sides or only when it costs me nothing?

Second, where am I tempted to perform? Most people do not lie outright. They perform. They adopt the expected cadence, the expected omissions, the expected enemy images. A stance fingerprint can reveal these performances because they create predictable alignment patterns with certain blocs. Seeing that pattern can interrupt unconscious compliance.

Third, where do I want principled friction? Some friction is a sign of incoherence. Other friction is a sign of integrity under pressure. A self-audit can help an author decide which blocs they are willing to disappoint and why. This is not about pleasing everyone. It is about being honest about the costs of one's commitments.

Used this way, the ten bits are not a branding weapon but a coherence instrument. They help an author avoid becoming a puppet of audience demands. They also help the author practice humility: recognizing where their own speech is ambiguous, opportunistic, or insufficiently disciplined.

The ethical advantage of this application is that it turns the classification inward. It treats the model as a tool for self-examination, where the primary risk is not harming others but confronting one's own inconsistencies.

10.4 AI-assisted sensemaking

AI systems can assist with stance-fingerprinting because they are good at pattern detection across large text corpora: summarizing recurring frames, identifying repeated moral grammar, and tracking shifts over time. But the same capacity makes misuse easier. If AI can generate fingerprints at scale, it can also generate labeling and targeting at scale. Therefore, AI-assisted use must be designed with accountability and constraint.

A responsible AI-assisted workflow has a clear division of labor.

AI can help with collection and organization: assembling a corpus, extracting long-form statements, grouping by topic, and producing neutral summaries of recurring claims. AI can also help with drift analysis: showing how language changes across time windows, highlighting key inflection points, and surfacing contradictions that deserve human review.

AI can help with candidate coding, but it should not be treated as the final judge. The safest use is to have the AI propose tentative bit settings with explicit uncertainty and a list of supporting excerpts, while the human makes the final determination. The human is responsible for context, intent inference restraint, and ethical application.

AI can also help with correlation awareness. It can warn the user when their chosen axes are not independent in practice, or when the coding appears to be following a coalition template. This is valuable because it resists the illusion that ten nominal bits provide ten independent degrees of identity.

Most importantly, AI can help enforce the humane protocol. It can be instructed to avoid naming private individuals, avoid generating targeting lists, require confidence scoring, require time windows, and require explicit statements of limitation. It can also be instructed to generate aggregate outputs rather than individual labels.

The central principle is accountability: AI can speed mapping, but it cannot replace judgment. Judgment includes moral caution, respect for persons, and the refusal to treat classification as destiny. A well-designed AI-assisted workflow makes it easier to see discourse pressure patterns while making it harder to turn those patterns into weapons.

Across all four applications, the best use of the 10-bit model is structural rather than personal. The model is most valuable when it reveals how public life forces people into binary performances—so that journalists, researchers, and authors can resist that coercion and restore space for reasoned, humane speech.

11. Conclusion

The “10-bit personality” model is intentionally modest. It does not claim to capture a person. It claims to capture a public signature: a repeatable pattern of alignment and friction between a speaker and ten audience blocs that impose binary demands. The reason this can feel powerful is that the public-figure regime is an unusual data environment. It penalizes silence, rewards legibility, and repeatedly forces statements into yes/no forms. Under those conditions, a code space of 2^10 = 1024 patterns can become practically discriminative within a reference set of

highly visible communicators, not because 1024 is large in the abstract, but because most people in that regime are fully observed and repeatedly interpreted.

Three forces make the model work at all. Missingness shrinks because public figures cannot remain un-coded on many axes for long; even silence is interpreted. Repeat exposure reduces ambiguity because a long archive of statements stabilizes patterns and makes meta-stance visible. Salience sharpens classification because the chosen axes correspond to high-conflict topics where organized audiences rapidly enforce boundaries. Together, these forces convert what looks like a weak combinatorial tool into a surprisingly informative map of cross-pressures.

But the same forces also limit the model. Ten bits often fail to identify individuals because public discourse does not inhabit the full code space. Coalitions bundle positions, causing correlation between bits and collapsing effective dimensionality. Media incentives and institutional pipelines produce templates, so many public figures become genre instances: mass-produced stance bundles with predictable fingerprints. Collisions follow naturally. Two people can share the same friction pattern while differing radically in reasons, virtues, temperament, and inner life. Even worse, audiences can generate friction through misunderstanding, strategic distortion, or decontextualized clips, producing fingerprints that reflect reception dynamics as much as speaker intent.

For these reasons, the safest interpretation of a stance fingerprint is structural, not personal. It is a tool for understanding how a discourse environment compresses people into legible codes, how loyalty tests operate, and how maximalist blocs demand rhetorical conformity. It is not a tool for declaring who someone “really is,” and it is not a substitute for reading, listening, or moral reasoning. The fingerprint can show where a speaker is being pulled, punished, or rewarded; it cannot tell you the speaker’s soul.

That is why the ethical burden is central, not optional. Compression tempts dehumanization because it makes persons portable as labels. If the model is used as a targeting device, a sorting mechanism, or a reputational weapon, it accelerates the very pathology it seeks to describe: the replacement of understanding with classification. A humane use requires provisional coding, explicit uncertainty, time-indexed snapshots, and a refusal to treat audience backlash as proof. Above all, it requires keeping the “human remainder” visible: the irreducible reality that a person exceeds their public stances, that reasons matter, that intent can diverge from reception, and that virtues and private faithfulness do not fit into a bitstring.

In the end, the 10-bit model is best understood as a mirror held up to the age: a description of how narrative war, coalition enforcement, and platform incentives press public life into binary forms. Its constructive value is not to help us label people faster, but to help us see the coercion

that makes labeling feel inevitable—and to practice speech, judgment, and restraint that restore dignity to persons in a world addicted to codes.

12. References

12.1 Polarization, Social Identity, Coalitional Pressure

- Fearon, J. D. (1994). Domestic political audiences and the escalation of international disputes. American Political Science Review, 88(3), 577–592. doi:10.2307/2944796
- Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., Mason, L., McGrath, M. C., Nyhan, B., Rand, D. G., Skitka, L. J., Tucker, J. A., Van Bavel, J. J., Wang, C. S., & Druckman, J. N. (2020). Political sectarianism in America. Science, 370(6516), 533–536. doi:10.1126/science.abe1715
- Huddy, L., Mason, L., & Aarøe, L. (2015). Expressive partisanship: Campaign involvement, political emotion, and partisan identity. American Political Science Review, 109(1), 1–17. doi:10.1017/S0003055415000604
- Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology: A social identity perspective on polarization. Public Opinion Quarterly, 76(3), 405–431. doi:10.1093/poq/nfs038
- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The origins and consequences of affective polarization in the United States. Annual Review of Political Science, 22, 129–146. doi:10.1146/annurev-polisci-051117-073034
- Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. Judgment and Decision Making, 8(4), 407–424.
- Kunda, Z. (1990). The case for motivated reasoning. Psychological Bulletin, 108(3), 480–498. doi:10.1037/0033-2909.108.3.480
- Mason, L. (2018). Uncivil Agreement: How Politics Became Our Identity. University of Chicago Press.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), The Social Psychology of Intergroup Relations (pp. 33–47). Brooks/Cole.
- Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. American Journal of Political Science, 50(3), 755–769.

12.2 Framing, Moral Psychology, and Speech Pressure

- Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. Journal of Communication, 43(4), 51–58. doi:10.1111/j.1460-2466.1993.tb01304.x

- Feinberg, M., & Willer, R. (2019). Moral reframing: A technique for effective and persuasive communication across political divides. Social and Personality Psychology Compass, 13(12), e12501. doi:10.1111/spc3.12501
- Goffman, E. (1974). Frame Analysis: An Essay on the Organization of Experience. Harper & Row.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. Journal of Personality and Social Psychology, 96(5), 1029–1046. doi:10.1037/a0015141
- Haidt, J. (2012). The Righteous Mind: Why Good People Are Divided by Politics and Religion. Pantheon.
- Noelle-Neumann, E. (1974). The spiral of silence: A theory of public opinion. Journal of Communication, 24(2), 43–51. doi:10.1111/j.1460-2466.1974.tb00367.x
- Sunstein, C. R. (2009). Going to Extremes: How Like Minds Unite and Divide. Oxford University Press. doi:10.1093/acprof:oso/9780195330436.001.0001
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. Science, 211(4481), 453–458. doi:10.1126/science.7455683

12.3 Misinformation, Conspiracy, and Network Dynamics

- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. Science, 348(6239), 1130–1132. doi:10.1126/science.aaa1160
- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. Proceedings of the National Academy of Sciences, 114(28), 7313–7318. doi:10.1073/pnas.1618923114
- Douglas, K. M., Sutton, R. M., & Cichocka, A. (2017). The psychology of conspiracy theories. Current Directions in Psychological Science, 26(6), 538–542. doi:10.1177/0963721417718261
- Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. Journal of Public Economics, 143, 51–68. doi:10.1016/j.jpubeco.2016.08.005
- Pennycook, G., & Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. Cognition, 188, 39–50. doi:10.1016/j.cognition.2018.06.011
- van Prooijen, J.-W. (2018). The psychology of conspiracy theories. Routledge.

- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. Science, 359(6380), 1146–1151. doi:10.1126/science.aap9559

12.4 Computational Methods for Discourse Mapping and Drift

- Blei, D. M. (2012). Probabilistic topic models. Communications of the ACM, 55(4), 77–84. doi:10.1145/2133806.2133826
- Blei, D. M., & Lafferty, J. D. (2006). Dynamic topic models. In Proceedings of the 23rd International Conference on Machine Learning (ICML). doi:10.1145/1143844.1143859
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of NAACL-HLT.
- Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. Journal of Economic Literature, 57(3), 535–574. doi:10.1257/jel.20181020
- Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016). Diachronic word embeddings reveal statistical laws of semantic change. In Proceedings of ACL.
- Küçük, D., & Can, F. (2020). Stance detection: A survey. ACM Computing Surveys, 53(1), Article 37. doi:10.1145/3369026
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., & Van Alstyne, M. (2009). Computational social science. Science, 323(5915), 721–723. doi:10.1126/science.1167742

12.5 Measurement, Coding Reliability, and Content Analysis

- Cohen, J. (1960). A coefficient of agreement for nominal scales. Educational and Psychological Measurement, 20(1), 37–46. doi:10.1177/001316446002000104
- Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. Psychological Bulletin, 76(5), 378–382. doi:10.1037/h0031619
- Krippendorff, K. (2011). Computing Krippendorff's alpha-reliability. Departmental Papers (ASC), University of Pennsylvania.
- Krippendorff, K. (2018). Content Analysis: An Introduction to Its Methodology (4th ed.). SAGE Publications.

12.6 Privacy, Reidentification, and Profiling Risks

- Dwork, C. (2006). Differential privacy. In Proceedings of ICALP. doi:10.1007/11787006_1

- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. Proceedings of the National Academy of Sciences, 110(15), 5802–5805. doi:10.1073/pnas.1218772110
- Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In Proceedings of the IEEE Symposium on Security and Privacy. doi:10.1109/SP.2008.33
- Sweeney, L. (2000). Simple demographics often identify people uniquely. Carnegie Mellon University, Data Privacy Working Paper 3.
- Youyou, W., Kosinski, M., & Stillwell, D. (2015). Computer-based personality judgments are more accurate than those made by humans. Proceedings of the National Academy of Sciences, 112(4), 1036–1040. doi:10.1073/pnas.1418680112

12.7 AI Governance, Human-AI Interaction, and Model Documentation

- Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. In Proceedings of CHI. doi:10.1145/3290605.3300233
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In Proceedings of FAccT. doi:10.1145/3442188.3445922
- Bostrom, N. (2014). Superintelligence: Paths, Dangers, Strategies. Oxford University Press.
- Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B., Anderson, H., Roff, H., Allen, G. C., Steinhardt, J., Flynn, C., Ó hÉigeartaigh, S., Beard, S., Belfield, H., Farquhar, S., Lyle, C., Crootof, R., & others. (2018). The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation. (Report).
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In Proceedings of FAT*. doi:10.1145/3287560.3287596
- National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). U.S. Department of Commerce.
- OECD. (2019). Recommendation of the Council on Artificial Intelligence. Organisation for Economic Co-operation and Development.

- Russell, S. (2019). Human Compatible: Artificial Intelligence and the Problem of Control. Viking.

12.8 Religion in Public Life and Pentecostal/Charismatic Practice

- Anderson, A. (2014). An Introduction to Pentecostalism: Global Charismatic Christianity (2nd ed.). Cambridge University Press.
- Casanova, J. (1994). Public Religions in the Modern World. University of Chicago Press.
- Habermas, J. (2006). Religion in the public sphere. European Journal of Philosophy, 14(1), 1–25. doi:10.1111/j.1468-0378.2006.00241.x
- Hollenweger, W. J. (1997). Pentecostalism: Origins and Developments Worldwide. Hendrickson Publishers.
- Rawls, J. (1993). Political Liberalism. Columbia University Press.
- Taylor, C. (2007). A Secular Age. Harvard University Press.
- The Holy Bible, New Testament. 1 Corinthians 12–14.

12.9 Nationalism, Populism, and Cross-National Audiences

- Anderson, B. (1983). Imagined Communities: Reflections on the Origin and Spread of Nationalism. Verso.
- Gellner, E. (1983). Nations and Nationalism. Cornell University Press.
- Mudde, C. (2007). Populist Radical Right Parties in Europe. Cambridge University Press.
- Müller, J.-W. (2016). What Is Populism? University of Pennsylvania Press.

12.10 Surveys and Reports Used as Context

- Pew Research Center. (2025, April 3). How the U.S. public and AI experts view artificial intelligence. (Report).
- Pew Research Center. (2025, September 17). How Americans view AI and its impact on people and society. (Report).
- Pew Research Center. (2025, October 15). How people around the world view AI. (Report).
- Pew Research Center. (2025, December 9). Teens, social media and AI chatbots 2025. (Report).

- Pew Research Center. (2025, October 3). How Americans view the Israel-Hamas conflict 2 years into the war. (Report).
- Pew Research Center. (2025, February 26). Decline of Christianity in the U.S. has slowed, may have leveled off. (Report).

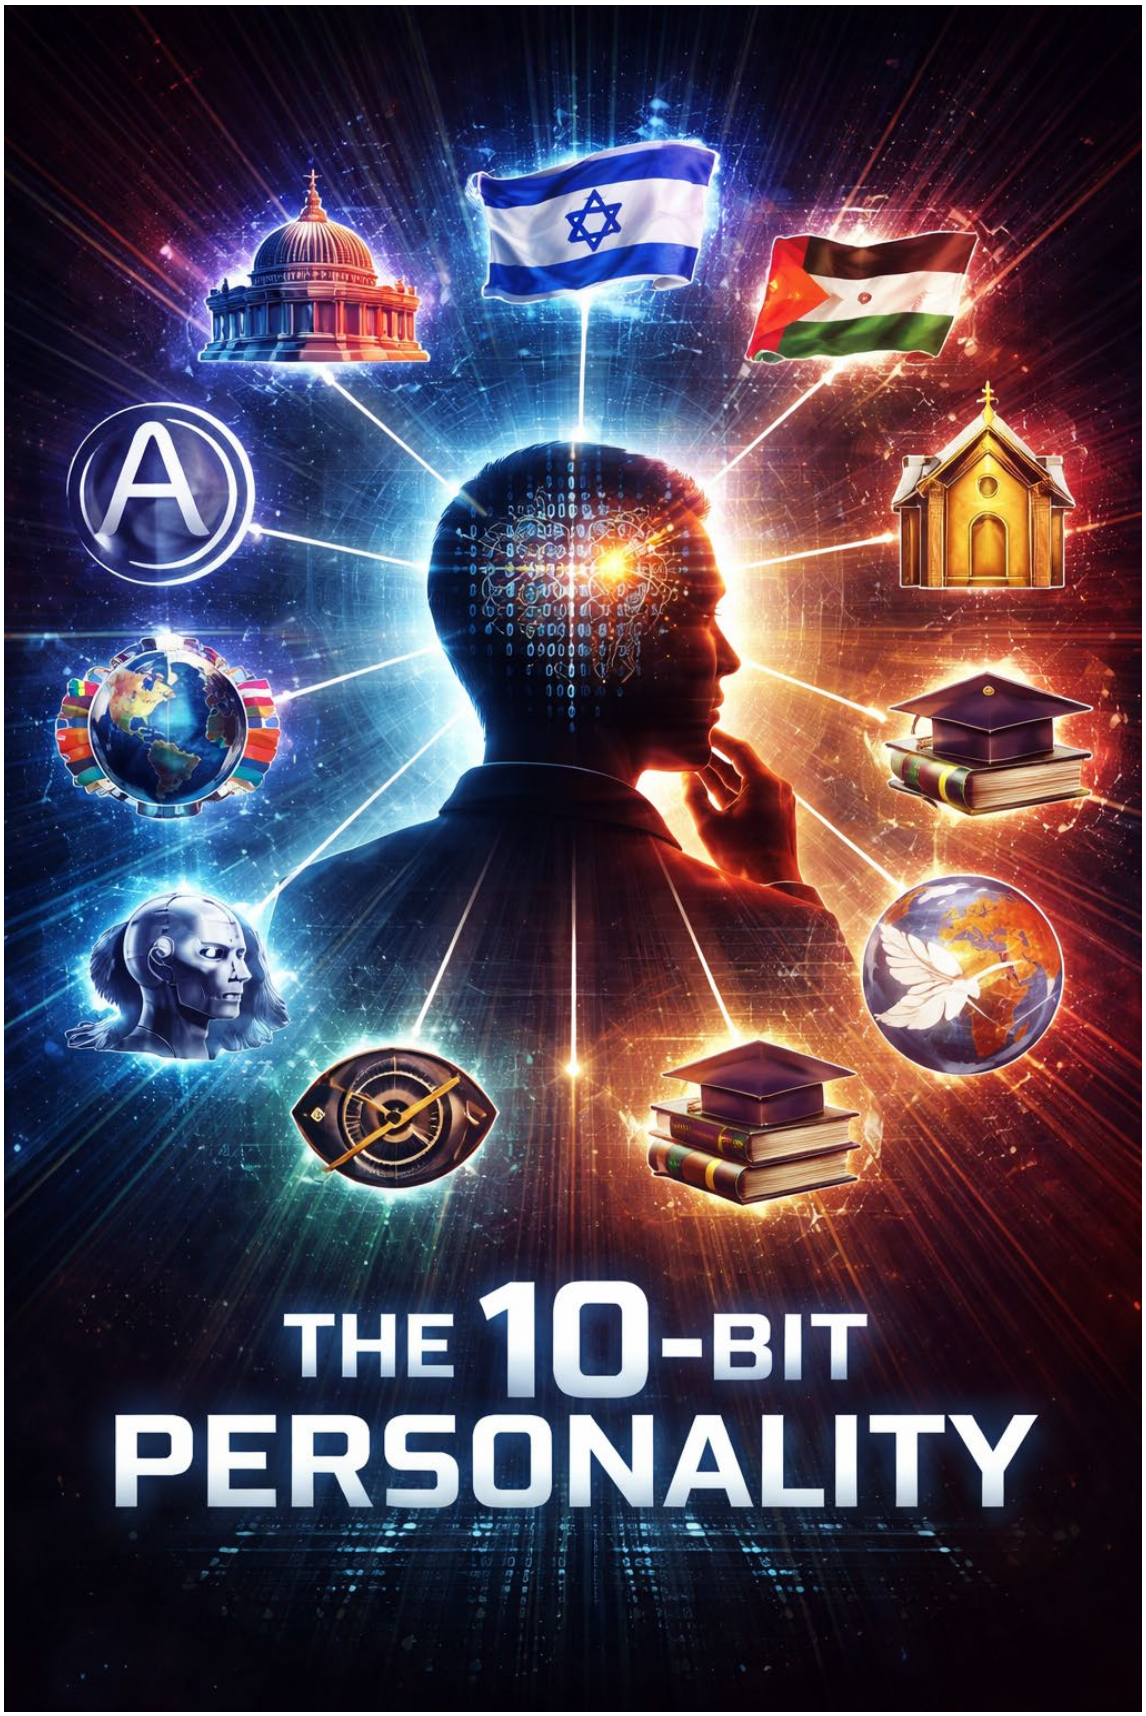
THE 10-BIT
PERSONALITY